Cross-convolutional-layer Pooling for Image Recognition

Lingqiao Liu, Chunhua Shen, Anton van den Hengel

Abstract—Recent studies have shown that a Deep Convolutional Neural Network (DCNN) trained on a large image dataset can be used as a universal image descriptor and that doing so leads to impressive performance for a variety of image recognition tasks. Most of these studies adopt activations from a single DCNN layer, usually the fully-connected layer, as the image representation. In this paper, we propose a novel way to extract image representations from two consecutive convolutional layers: One layer is used for local feature extraction and the other layer serves as guidance to pool the extracted features. By taking different viewpoints of convolutional layers, we further develop two schemes to realize this idea. The first one directly uses convolutional layers from a DCNN. The second one applies the pre-trained CNN on densely sampled image regions and treats the fully-connected activations of each image region as convolutional layer feature activations. We then train another convolutional layer on top of that as the pooling-guidance convolutional layer. By applying our method to three popular visual classification tasks, we observe that our first scheme tends to perform better on the applications that need strong discrimination on subtle object patterns within small regions while the latter excels in the cases that require discrimination on category-level patterns. Overall, our proposed method achieves superior performance over existing ways of extracting image representations from a DCNN. In addition, we apply cross-layer pooling to image retrieval and propose methods to reduce the computational cost. Experimental results suggest that the proposed method also leads to considerable performance improvement on the task of image retrieval.

Index Terms—Convolutional networks, deep learning, pooling, fine-grained object recognition.

arXiv:1510.00921v2 [cs.CV] 9 Jun 2016

1 Introduction

Recently, Deep Convolutional Neural Networks (DCNNs) have attracted much attention in visual recognition, largely due to their excellent performance [1]. It has been discovered that activations of a DCNN trained on a large dataset, such as ImageNet [2], can be employed as a universal image descriptor, and applying this descriptor to many visual classification and retrieval problems delivers impressive performance [3], [4], [5]. This discovery quickly sparked significant research interest, and inspired a number of extensions, including [6], [7]. A fundamental issue with these methods is how to generate an image representation from a pre-trained DCNN. Most current solutions take activations of a single DCNN layer, often the fully-connected layer, as the image representation.

In this paper, we show that we can build a powerful image representation from the activations extracted from two consecutive convolutional layers. We term our method as cross-convolutional layer pooling (or cross-layer pooling for short). This new method relies on two crucial components. (a) We extract local features from one convolutional layer; and (b) we pool extracted local features by using activations from its successive convolutional layer as guidance.

The first component is motivated by recent works [6], [7], [8] which have shown that DCNN activations are not translation invariant and that it is beneficial to extract fully connected layer activations from a DCNN to describe local regions and create the image representation by pooling multiple regional DCNN activations. In this paper, we view those regional CNN activations as a newly added convolutional layer (named as the augmented convolutional layer as discussed in Section 3.1). Inspired by this view, we also extract local features from the original convolutional layers of the pre-trained CNN.

The second component is motivated by the parts-based pooling method [9] which was originally proposed for fine-grained image classification. This method creates one pooling channel for each detected part region and the final image representation is obtained by concatenating pooling results from the multiple channels. We generalize this idea to the context of DCNNs and avoid the need for annotating predefined parts. More specifically, we deem the feature map of each filter in a convolutional layer as the detection score map of a part detector and apply the feature map to weight regional descriptors extracted from the previous convolutional layer in the pooling process. The final image representation is then computed by concatenating pooling results from multiple channels with each channel corresponding to one feature map. Note that unlike existing regional-DCNN based methods [6], [7], the proposed method does not need any additional dictionary learning or encoding steps at both training and testing stages once the convolutional layer activations is available. To further reduce the memory usage in storing image representations, we also experiment with a coarse 'feature sign quantization' compression scheme and show that the discriminative power of the proposed representation can be largely maintained after compression.

Moreover, we explore the use of cross-layer pooling for image retrieval. To cope with the expensive computational cost of the direct implementation of cross-layer pooling, we propose to employ feature binarization and adaptive

• *The authors are with School of Computer Science, The University of Adelaide, Adelaide, SA 5005, Australia (E-mail: {lingqiao.liu, chunhua.shen, anton.vandenhengel}@adelaide.edu.au).*

selection of pooling channels to significantly reduce the computational cost.

We conduct extensive experiments on three popular visual classification datasets and three popular image retrieval datasets for both tasks of classification and retrieval. Experimental results show that the proposed method can achieve significantly better performance than competitive methods in most cases. We carry out further ablation studies to analyze the importance of various components of our approach.

A preliminary version of this paper appeared in [10]. In this work, we have made the following extensions.

- We view the scheme of extracting fully connected CNN activations at densely sampled regions as a newly added convolutional layer and perform cross-layer pooling at that level. This extension makes our method more widely applicable.
- We apply cross-layer pooling to a new task of image retrieval and propose new schemes to reduce the computational cost.
- We have presented more experiments, including experiments with a better CNN model and new ablation studies in this paper.

Preliminaries. Two network structures are considered in this paper. One is the AlexNet [1] and another is the VGG very deep network (VGGNet) [11]. Both networks are composed of a few convolutional layers and fully connected layers. We use the notion fc-k and conv-k to denote the k-th fully connected layer and k-th convolutional layer respectively. At each convolutional layer, multiple filters are applied, and it results in multiple feature maps, one for each filter. In this paper, we use the term 'feature map' to indicate the convolutional result (after applying the ReLU) of one filter and the term 'convolutional layer activations' to indicate feature maps of all filters in a convolutional layer.

2 Related work

In the literature, there are two main strategies of using a pre-trained DCNN to create image representations: (a) directly feeding the whole image into a pre-trained DCNN and extracting its activations as the image representation; (b) applying the pre-trained DCNN to the subregions of the input image and aggregating activations from multiple regions as the image representation. Usually, the first strategy extracts the last few fully-connected layers' activations as the image-level representation. Fine-tuning can be applied to make the network better adapted to a target task. Furthermore, to make the method more robust to image transforms, averaging activations from several jittered versions of the original image, e.g., several slightly shifted versions of the input image, can be employed to achieve improved classification performance [4].

DCNNs can also be applied to extract local features. It is suggested that DCNN activations are not invariant to a large amount of translation [6] and the performance can be degraded if input images are not well aligned. To make the feature more robust to translation, a simple idea is to sample multiple regions from an input image and use one DCNN, called regional-DCNN in this scenario, to describe each region. The final image representation is aggregated from activations of those regional-DCNNs [6]. In [6], another layer of unsupervised encoding is employed to create the image-level representation [6], [7]. In [8], discriminative patterns are mined from those regional activations for classification. It is shown that for many computer vision tasks [6], [7], region-DCNNs methods often lead to improved performance than directly extracting DCNN activations as global features.

One common factor in the above methods is that almost all of them use fully-connected layer activations as features. Some recent studies explore the use of convolutional layers to extract image representations. For example, the 'Deep Texture' work [12] applies Fisher vector pooling to the local features extracted from a convolutional layer to create image representations for texture classification. The work of [13] use the pooled convolutional activations for object detection. In [5], the authors discover that the pooled convolutional features work very well on image retrieval tasks. The work in [14] is most relevant to our work. As mentioned in the paper of [14], their work is an extension of the proposed method here (our conference version [10]) for fine-grained image classification. It uses a similar strategy as ours to combine the convolutional feature activations from two DCNNs and jointly fine-tune all the parameters in an end-to-end fashion.

3 Our proposed method

An overview of the proposed method is shown in Figure 1.

3.1 Convolutional layers vs. fully connected layers

One major difference between convolutional and fully-connected layer activations is that the former is embedded with rich spatial structure while the latter is not. Convolutional layer activations can be formulated as a tensor of the size $H \times W \times D$, where H, W denote the height and width of each feature map and D denotes the number of feature maps. Essentially, the convolutional layer divides the input image into $H \times W$ regions and uses D-dimensional filter responses to describe the visual pattern within each region. Thus, convolutional layer activations can be viewed as a 2-D array of D-dimensional *local features* with each one describing a local region. For the sake of clarity, we term each of the $H \times W$ regions as a *spatial unit*, and the D-dimensional feature maps corresponding to a spatial unit as the *feature vector in a spatial unit*. A fully-connected layer can take a convolutional layer's activations as the network input and transforms them into a feature vector representing the whole image. Spatial structure is lost through this process, meaning that the feature vector corresponding to a particular spatial area cannot be recovered from the activations of the subsequent fully-connected layer.

As mentioned in Section 2, DCNNs can also be applied to image patches to extract local features from the fully-connected layer, as a means of compensating the fact that they are not translation invariant. Note that if the regions are densely sampled from the input image and they are of the same size, the above procedure can be approximately

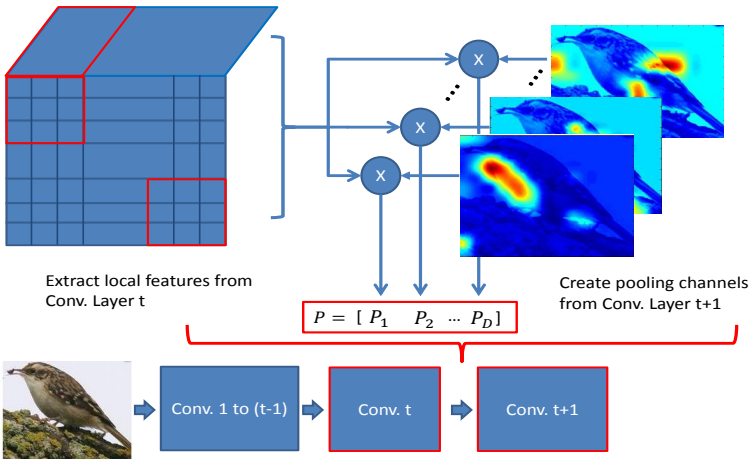

Fig. 1. Overview of the proposed method. Our method performs the pooling operation by using two consective convolutional layers.

achieved[1] by firstly resizing the image into a larger size to extract activations of the last convolutional layer and then stacking another convolutional layer with the original fully-connected layer coefficients as the filters. In other words, this strategy turns the fully-connected layer into a convolutional layer. For example, in the AlexNet, the fc-6 layer maps activations of conv-5 (after max-pooling), which is the size of 6×6×256, into a 4096-dimensional vector. fc-6 can also be seen as a filter bank with the support region of 6×6 spatial units, 256 input channels, and 4096 filters. Thus, to extract fc-6 activations from regions with 1/2 size of the height and width of the original image can be implemented by applying the following two steps: (a) resizing the input image by a factor of two, which creates conv-5 (after max-pooling) activations with the size of 12×12×256; and (b) performing the convolutional operation with a 6×6×256×4096 kernel.

For the clarity of reference, we term this type of convolutional layer as the 'Augmented Convolutional Layer' or the 'AConv' layer in short and the original convolutional layer in the DCNN as the 'OConv' layer.

3.2 Cross-convolutional-layer Pooling

In [6], [7], it has been shown that applying an additional pooling operation on the local features extracted from multiple image regions can significantly boost classification performance. Motivated by these methods, we can design a specific pooling layer and apply it to the local features extracted from a convolutional layer which can be either the OConv layer or the AConv layer.

Various pooling methods may be chosen to aggregate the local features, e.g., max-pooling, sum-pooling or Fisher vector based pooling [15]. In this section, we propose an alternative pooling method which can significantly improve classification performance.

The proposed method is inspired by the parts-based pooling strategy [9] mainly used in fine-grained image classification. There, multiple regions-of-interest (ROI) are first detected, with each corresponding to one human-specified object part, e.g., the tails of birds. Then local features falling into each ROI are then pooled together to obtain a pooled feature vector. Given D object parts, this strategy creates D different pooled feature vectors and these vectors are concatenated together to form the final image-level representation. It has been shown that this simple strategy achieves significantly better performance than directly pooling all local features together. Formally, the pooled feature from the kth ROI, denoted as $\mathbf{P}_k^t$, can be calculated by the following equation (we consider sum-pooling in this case):

$$\mathbf{P}_k^t = \sum_{i=1} \mathbf{x}_i I_{i,k}, \tag{1}$$

where $\mathbf{x}_i$ denotes the ith local feature and $I_{i,k}$ is a binary *indicator map* indicating whether $\mathbf{x}_i$ falls into the kth ROI. We can also generalize $I_{i,k}$ to real values with its value indicating the 'membership' of a local feature to a ROI. Essentially, each indicator map defines a pooling channel and the image representation is the concatenation of pooling results from multiple channels.

However, in a general image classification task, usually human-specified parts annotation is not available. Even worse, for many fine-grained image classification tasks, annotation and detection of these parts can be nontrivial or impossible. To solve this difficulty, in this paper we propose to use feature maps of the $(t+1)$th convolutional layer as D_{t+1} indicator maps. By doing so, D_{t+1} pooling channels are created for the local features extracted from the t-th convolutional layer. We name this method *cross-convolutional-layer pooling* or *cross-layer pooling* in short. The use of feature maps as indicator maps is motivated by the observation that a feature map of a deep convolutional layer is usually sparse and indicates some semantically meaningful regions[2]. This observation is illustrated in Figure 2. In Figure 2, we choose three images taken from the Birds-200 [17] dataset. We sample three feature maps from 256 feature maps in conv5 and overlay them on the original images for better visualization. As can be seen from Figure 2, the activated regions of the sampled feature map (highlighted in warm color) are semantically meaningful. For example, the activated region in the first row tends to localize at the head region of a bird while the feature map shown in

1. Slight difference might be introduced due to the padding operations.

2. Note that similar observation has also been made in [16].

the second row exhibits high values around the claw area. Thus, the filter of a convolutional layer works as a part detector, and its feature map serves a similar role as the part region indicator map. Certainly, compared with the parts detector learned from human-specified part annotations, the filter derived from a convolutional layer may not be directly task-relevant. However, the discriminative power of our image representation can benefit from combining a much larger number of indicator maps, e.g., 256 as opposed to 20∼30 (the number of parts that are manually annotated in fine-grained image classification datasets), which is akin to applying bagging to boost the performance of multiple weak classifiers.

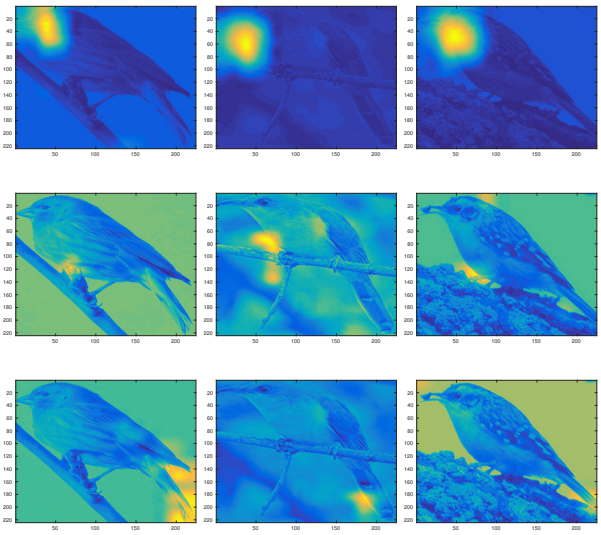

Fig. 2. Visualization of some feature maps extracted from the 16th convolutional layer of the VGGNet. Three feature maps and their activations on three different images are shown. Each row represents the feature map corresponding to the same filter. Warmer color indicates higher activation value.

Formally, the image representation extracted from cross-layer pooling can be expressed as follows:

$$\mathbf{P}^t = [{\mathbf{P}_1^t}^\top, {\mathbf{P}_2^t}^\top, \cdots, {\mathbf{P}_k^t}^\top, \cdots, {\mathbf{P}^t}_{D_{t+1}}^\top]^\top$$
$$\text{where,} \quad \mathbf{P}_k^t = \sum_{i=1}^{N_t} \mathbf{x}_i^t a_{i,k}^{t+1}. \tag{2}$$

Here $\mathbf{P}^t$ denotes the pooled feature for the t-th convolutional layer, which is calculated by concatenating the pooled feature of each pooling channel $\mathbf{P}_k^t, k = 1, \cdots, D_{t+1}$. $\mathbf{x}_i^t$ denotes the i-th local feature in the t-th convolutional layer. $a_{i,k}^{t+1}$ is the activation value of the i-th spatial unit and the k-th filter at the $(t+1)$-th convolutional layer. Here we assume that there is a correspondence between the i-th local feature at the t-th layer $\mathbf{x}_i^t$ and the i-th feature activations at the $(t+1)$-th layer. This correspondence can be easily established if the consective convolutional layers have the same spatial unit layout. For example, the last two convolutional layers in the Alex net both have a 13×13 spatial unit layout and we can deem that the feature maps at the same spatial unit across two layers are corresponded.

Fig. 3. This figure demonstrates the image style mismatch issue when using fully-connected layer activations as regional descriptors. Top row: input images that a DCNN 'sees' at the training stage. Bottom row: input images that a DCNN 'sees' at the test stage.

3.3 Augmented convolutional layer vs. original convolutional layer

To implement cross-layer pooling, one needs to specify two convolutional layers. In practice, these two convolutional layers can either be chosen from the AConv layer or the OConv layer[3]. A question to answer is, which type of convolutional layers perform better? Our discovery suggests that the best performed convolutional layer type can be different for different recognition tasks.

For an AConv layer, its convolutional filters are the fully-connected layer parameters of the original CNN. Therefore an AConv layer encodes higher-level visual concepts information, such as object-category-level information. Thus, if the target problem involves identification of some object-category-level patterns, e.g., to classify if a "car" occurs in the image, then an AConv layer should be used. In this case, its activations work similarly as the object bank detectors [18]. Note that even if the problem does not directly involve the identification of an object that appears in the network training task, e.g., the 1000 categories in the ImageNet dataset, the category-level pattern detection may still be beneficial. For example, for scene classification, the occurrence of an object, e.g., 'bed lamp', can be a strong indicator of the scene class 'bedroom'.

For some other applications, it is often required to discriminate subtle patterns at relatively small key regions, such as the fine-grained image classification. If we extract features from an AConv layer to capture those key regions, we may need a smaller receptive field to make the DCNN focus on the patterns within these regions. However, when the receptive field of the AConv layer becomes smaller, the input of the network tends to become very different to the input (the training images) of the network at its training stage.

Figure 3 shows an example of this case. The top row of Figure 3 shows some images from the ImageNet dataset which is commonly used to train the DCNN while the bottom row shows some input regions to the AConv layer (resized to the same image size as the top row) on the Birds-200 [17] dataset. As can be seen, although they all have the same image size, their appearance and level of detail are very different. Thus, applying an AConv layer and extracting fully-connected layer activations as local

3. It is also possible to choose one OConv layer and one AConv layer to perform cross-layer pooling. We discuss this possibility in Section 4.2.1.

features introduce a significant input image style mismatch, which can potentially undermine the discriminative power of DCNN activations. In contrast, directly using the original convolutional feature activations can mitigate the above drawback. This is because to extract all the local features from the original convolutional layer need to apply *a single DCNN forward computation to a whole image once*. Consequently, the input of a DCNN is still a whole image but not a small local patch. This alleviates the domain mismatch issue when we use the AConv layer.

3.4 Implementation details

PCA: In our implementation, we perform PCA on $\mathbf{x}_i^t$ to reduce the dimensionality of $\mathbf{P}^t$ to 2000 dimensions for the AConv layer. For the OConv layer, we still perform PCA to de-correlate the local feature but without performing dimensionality reduction. Empirically we find that this leads to slightly improved performance than using the original OConv local features.

Normalization: Since the number of activated spatial units at the guidance convolutional layer can be different for different pooling channels. The pooling vector derived from different channels may contain very different magnitude of energies. Thus, in our implementation we ℓ_2 normalize the pooled coding vector for each channel; and then we apply power normalization to $\mathbf{P}^t$. That is, we use $\hat{\mathbf{P}}^t = \text{sign}(\mathbf{P}^t)\sqrt{|\mathbf{P}^t|}$ as the image representation to further improve performance.

Feature sign quantization: Besides the aforementioned processing to create the image representation. We have also tried directly using $\text{sign}(\mathbf{P}^t)$ as an image representation. That is, we coarsely quantize $\mathbf{P}^t$ into $\{-1, 1, 0\}$ according to the feature sign of $\mathbf{P}^t$.

Adding a new convolutional layer for the AConv layer: One issue when using the AConv layer for cross layer pooling is that we need to find two consecutive AConv layers. These two layers can be obtained by using two consecutive fully-connected layers in the original DCNN. However, since the fully-connected layers in most commonly used DCNN models have a very large number of output neurons, e.g., 4096 or 1000. Directly performing cross-layer pooling on those AConv layers would result in a very high dimensional image representation. To solve this issue, we only use one fully connected layer from the original DCNN as one AConv layer, and stack another newly added convolutional layers on top of it with a much smaller number of filters, e.g., 100.

Then we train the new convolutional layer on the target dataset. The network architecture of our implementation is as follows: a max-pooling layer is applied to pool the activations of the newly added convolutional layer and the pooled result is feed into a logistic regression layer. The negative entropy loss is then used to train the new convolutional layer.

3.5 Application to image retrieval

Recently, it has been discovered that the pooled convolutional layer activations can be a good image representation for image retrieval [5]. Observing the success of the work in [5], we apply our cross-layer pooling method to image retrieval. Since cross-layer pooling creates multiple pooling channels with each pooling channel capturing one type of visual pattern and the pooled result of each channel is normalized, image representations created by cross-layer pooling can depict various aspects of visual patterns within the image in a balanced form. In comparison, information in the representation generated by the direct pooling method in [5] may be dominated by patterns that occur most frequently within the given image.

However, directly applying cross-layer pooling to the task of image retrieval would incur a very high computational cost due to the high dimensionality of the generated image representations. For example, if there are M pooling channels, the computational cost can be M times higher than the method in [5]. To alleviate this drawback, here we propose two strategies. The first strategy is to simply binarize the cross-layer pooling feature. This is inspired by the observation (which will be experimentally demonstrated in Section 4.4) that keeping the sign of feature values[4] are unlikely to affect much on the discriminative power of cross-layer pooling.

The second strategy is to adaptively select a few, e.g., k pooling channels for *each query image* and only compare the feature similarity on the feature pooled from the selected channels. Formally, for such a scheme the image similarity between a query image and a test image is calculated via:

$$s(\mathbf{I}_q, \mathbf{I}) = \sum_{k \in \mathcal{S}} \mathbf{q}_k^T \mathbf{x}_k, \tag{3}$$

where $\mathbf{I}_q$ and $\mathbf{I}$ are the query image and a test image. In original cross-layer pooling, both the query image and the test image are represented by D subvectors which are pooled from each pooling channel created by a convolutional layer. We use $\mathbf{q}_k$ and $\mathbf{x}_k$ denote the kth subvectors of the query and a test image respectively. For the original cross-layer pooling approach, comparison should be made on all D subvectors. Here in Equation (3), only subvectors whose indices fall into a small subset $\mathcal{S}$ ($|\mathcal{S}| = k \ll D$) are compared. Thus, the computational cost can be greatly reduced in comparison with the naive implementation of cross-layer pooling for image retrieval. In this paper, we construct the set $\mathcal{S}$ by selecting channels (feature maps of a convolutional layer) with top k maximal average activations. By applying this criterion, the convolutional feature maps with less significant activations will be discarded. Thus, for this operation one additional benefit besides reduction of the computational cost is that it might suppress the noise patterns and therefore potentially improve retrieval performance. Note that besides selecting channels with top k largest activation values, other criteria can be applied. For example, if the retrieval task is to find a specific object type, e.g., sculpture, cloth, an importance weight for each pooling channel can be learned by using images which contain the object-of-interest as positive training samples and irrelevant images as negative training samples.

4. As will be discussed in the second strategy, some channels can be discarded during feature similarity calculation, which is equivalent to setting the feature values within the discarded channels to 0. In this view, there are still three possible values, $1, -1, 0$, in the resulted image representation.

4 EXPERIMENTS

We organize our experiments into three parts, the first part evaluates the proposed cross-layer pooling method for the application of image classification. Then in the second part, we conduct ablation study and demonstrate the impact of various components of our method. In the third part, we evaluate the proposed method on image retrieval tasks. The focus of the third part is to evaluate if the proposed cross-layer pooling can lead to better performance than the recent method in [5] which also uses a convolutional layer pooling strategy for image retrieval.

4.1 Experiments for image classification

4.1.1 Experimental protocol

TABLE 1
Comparison of results on MIT-67. The lower part of this table lists some results reported in the literature.

Methods	Accuracy	Remark/Setting
CNN-Global	57.9%	AlexNet
CNN-Jitter	61.1%	AlexNet
SCFV [7]	59.2%	AlexNet, OConv
CrossLayer (proposed)	63.0%	AlexNet, OConv
SCFV [7]	**68.2%**	AlexNet, AConv
CrossLayer (proposed)	**68.2%**	AlexNet, AConv
CNN-Global	68.2%	VGGNet
CNN-Jitter	70.2%	VGGNet
SCFV [7]	73.5%	VGGNet, OConv
CrossLayer (proposed)	**74.4%**	VGGNet, OConv
SCFV [7]	76.4%	VGGNet, AConv
CrossLayer (proposed)	**78.2%**	VGGNet, AConv
Fine-tuning [4]	66.0%	fine-tunning on MIT-67
MOP-CNN [6]	68.9%	three scales
VLAD level2 [6]	65.5%	single scale
CNN-SVM [3]	58.4%	-
FV+DMS [19]	63.2%	-
DPM [20]	37.6%	-
DeepTexture [12]	**81.7%**	7 scales
Texture Synthesis [21]	75.0%	using the Gram matrix on fc18 layer (VGGNet) for classification

We evaluate the proposed method on three datasets: MIT indoor scene-67 (MIT-67) [22], Caltech-UCSD Birds-200-2011 [17] (Birds-200) and PASCAL VOC 2007 [23] (PASCAL-07) for image classification. These three datasets cover several popular topics in image classification, namely, scene classification, fine-grained object classification and generic object classification.

We compare the proposed method against three baselines, which are:
(a) directly using fully-connected layer activations for the whole image (CNN-Global);
(b) averaging fully-connected layer activations from several transformed versions of an input image. Following [3], [4], we transform the input image by cropping its four corners and middle regions as well as by creating their mirrored versions;
(c) applying the sparse coding based Fisher vector encoding method [7] on the local feature extracted from the convolutional layer (in both schemes).

Since the SCFV method has demonstrated superior performance to the MOP method in [6], we do not include MOP in our comparison. To make the comparison fair, we reimplement all three baseline methods. Also, If applicable, the same post-processing steps such as PCA, ℓ_2 normalization and power normalization, are applied to the baseline methods.

Two CNN models are adopted throughout our experiment: the AlexNet [1] and the VGG very-deep 19 layers network [11]. Two different types of convolutional layers are used, the original convolutional layer (denoted as OConv) and the augmented convolutional layer (denoted as AConv). As discussed in Section 3.1, the latter can be implemented by applying DCNN on a set of image regions which are extracted on a dense grid. In our implementation, we first resize the input image to 512×512 pixels and then extract image regions with the size 128×128 at a step size of 32 pixels. For OConv layers, we report the results obtained by using the 4-th and 5-th convolutional layers for the AlexNet and the 15-th and 16-th convolutional layers for the VGGNet since those settings achieve the best performance. We also explore the use of other convolutional layers in the second part of our experiment. For the AConv layer, we extract local features from the fc6 and fc18 in the AlexNet and the VGGNet respectively. Then we stack a new convolutional layer with 100 filters on top of them and train the newly added layer on the target dataset. This new layer is trained with a learning rate of 0.01 by 50 epochs. No data augmentation is used in this training step.

We use Libsvm [24] as the SVM solver and use precomputed linear kernels as inputs. This is because the calculation of linear kernels/Gram matrices can be easily implemented in parallel. When feature dimensionality is high the kernel matrix computation actually dominates most of computational time. Thus it is appropriate to use parallel computing to accelerate this process.

4.1.2 Classification results

Scene classification: MIT-67. MIT-67 is a commonly used benchmark dataset for evaluating scene classification methods. It contains 6700 images with 67 indoor scene categories. Following the standard setting, we use 80 images in each category for training and 20 images for testing. The results are shown in Table 1. It can be seen that the proposed cross-layer pooling achieves the overall best performance in most settings. The best performance is achieved by using cross-layer pooling and the AConv layer: this setting produces 68.2% classification accuracy for the AlexNet and 78.2% for the VGGNet.

Furthermore, it is clear that extracting local features from the AConv layer, as has been done in SCFV and CrossLayer, achieves significant performance improvement in comparison with global CNN features, i.e., Global and Jitter. Finally, the use of the VGGNet further boost the classification performance by a large margin.

By comparing with the performance reported from the literature, we can see that the proposed method surpasses most state-of-the-art methods. The only exception is the result in [12]. However, its method is very close to our SCFV (with OConv) baseline and its good performance is largely due to their brute force mutiple-scale strategy. They have utilized 7 scales while we only use one single scale.

Fine-grained image classification: Birds-200. Birds-200 is the most popular dataset for fine-grained image classifica-

TABLE 2
Comparison of results on Birds-200. Note that the method annotated with "use parts" requires parts annotations and detection while our methods do not employ these annotations so they are not directly comparable with our method.

Methods	Accuracy	Remark
CNN-Global	59.2%	no parts. AlexNet
CNN-Jitter	60.5%	no parts. AlexNet
SCFV [7]	64.2%	no parts, AlexNet, OConv
CrossLayer	**73.3%**	no parts, AlexNet, OConv
SCFV [7]	66.4%	no parts, AlexNet, AConv
CrossLayer	71.7%	no parts, AlexNet, AConv
CNN-Global	62.5%	no parts. VGGNet
CNN-Jitter	63.6%	no parts. VGGNet
SCFV [7]	73.7%	no parts, VGGNet, OConv
CrossLayer	**77.0%**	no parts, VGGNet, OConv
SCFV [7]	66.2%	no parts, VGGNet, AConv
CrossLayer	69.4%	no parts, VGGNet, AConv
GlobalCNN-FT [4]	66.4 %	no parts, fine tunning
Parts-RCNN-FT [25]	76.37 %	use parts, fine tunning
Parts-RCNN [25]	68.7 %	use parts, no fine tunning
CNNaug-SVM [3]	61.8%	-
CNN-SVM [3]	53.3%	CNN global
DPD+CNN [26]	65.0%	use parts
DPD [27]	51.0%	-
Bilinear CNN [27]	77.9%	Two networks
Bilinear CNN [27]	**81.9%**	Two networks, fine-tuning
Texture Synthesis [21]	67.3	using the Gram matrix on conv16 layer (VGGNet) for classification

tion. It contains 11788 images with 200 different bird species. This dataset provides ground-truth annotations of bounding boxes and parts of birds, e.g. the head and the tail, on both the training set and the test set. In this experiment, we just use the bounding box annotation. The results are shown in Table 2. As can be seen, the cross-layer pooling achieves the best classification performance: 77.0% when the VGGNet model is used.

Also, using the original convolutional layer achieves much better performance than the use of the AConv layers. For both the AlexNet and the VGGNet models, the best performance is achieved by using features from the OConv layer. The underlying reason can be well explained in Section 3.3. That is, the discriminative information of birds species usually lies at small regions and it is more appropriate to extract features from original convolutional layers due to the image style mismatch issue discussed in Section 3.3.

Our best performance is among the best on this dataset. The work in [14] reports higher classification accuracy than ours. However, it relies on a fine-tuning step on two different networks and it adopts some different experimental settings. Namely, their convolutional layers have different number of spatial units to us[5] (28×28 as oppose to 14×14 in our experiment); it performs decision score calibration on the SVM while we only use the standard one-vs-rest SVM.

Object classification: PASCAL-2007. PASCAL VOC 2007 has 9963 images of 20 object categories. The task is to predict the presence of each object in each image. Note that most object categories in PASCAL-2007 are also included in ImageNet which is the training set of the AlexNet and VGGNet models. ImageNet can be seen as a super-set of PASCAL-2007. The results on this dataset are shown in Table 3.

5. When the same spatial units configuration is used, our cross-layer pooling can achieve 80% classification accuracy which is closer to the result in [14]. Note that we only use single network and does not apply fine-tuning and score calibration.

TABLE 3
Comparison of results on PASCAL VOC 2007.

Methods	mAP	Remark
CNN-Global	71.7%	AlexNet
CNN-Jitter	75.0%	AlexNet
SCFV [7]	66.8%	AlexNet, OConv
CrossLayer	71.3%	AlexNet, OConv
SCFV [7]	76.9%	AlexNet, AConv
CrossLayer	**79.1%**	AlexNet, AConv
CNN-Global	83.6%	VGGNet
CNN-Jitter	84.5%	VGGNet
SCFV [7]	82.9%	VGGNet, OConv
CrossLayer	84.1%	VGGNet, OConv
SCFV [7]	85.1%	VGGNet, AConv
CrossLayer	**87.1%**	VGGNet, AConv
CNNaug-SVM [3]	77.2%	with augmented data
CNN-SVM [3]	73.9%	no augmented data
NUS [28]	70.5%	-
GHM [29]	64.7%	-
AGS [30]	71.1%	-
Texture Synthesis [21]	84.7%	using the Gram matrix on fc18 layer (VGGNet) for classification

From Table 3, we can see that again the best performance is achieved by using cross-layer pooling and the VGGNet model. Unsurprisingly, the AConv layer performs better than the OConv layer in this dataset because the training categories of the DCNN overlaps with PASCAL-2007 and the AConv layer contains this category-level information. The per-class performance of three best comparing methods, namely, CNN jitter with VGGNet, SCFV with the AConv layer from VGGNet and cross-layer pooling with the AConv layer from VGGNet, is shown in Table 4. As can be seen, the proposed cross-layer pooling achieves the best performance in most classes.

4.2 Ablation study

From the above experiments, the advantage of using the proposed method has been clearly demonstrated. In this section we further examine the effect of various components in our method.

4.2.1 Using different convolutional layers

First, we are interested in examining the performance of using convolutional layers other than the 4th and 5th convolutional layers in the AlexNet and the 15-th and 16-th convolutional layers in the VGGNet. We investigate the performance of using the 3th and 4th convolutional layers for the AlexNet and the 14th and 15th convolutional layers in VGGNet. The results are shown in Table 6. From the results we can see that using 4-5th layers and 15-16th layers achieves superior performance over using the 3-4th layers and 15-16th layers. This is coincident with the existing discovery [16] that the deeper the convolutional layer, the better discriminative power it has.

As discussed in Section 3.1, the process of extracting features from a fully-connected layer at multiple local regions can be viewed as applying a special convolutional

TABLE 4
Comparison of results on Pascal VOC 2007 for each of 20 classes. The classes on which cross-layer pooling achieves significant better performance are labeled black.

	TV	train	sofa	sheep	plant	person	mbike	horse	dog	table
Global Jitter (VGGNet)	81.9	96.3	72.9	86.1	61.6	95.2	89.3	91.5	90.9	79.7
SCFV (VGGNet)	82.3	95.7	78.9	84.0	62.6	95.8	89.1	94.1	90.4	82.2
CrossLayer (VGGNet)	**85.3**	96.7	**80.1**	**87.6**	**64.6**	**96.7**	**91.7**	94.3	**93.0**	82.1
	cow	chair	cat	car	bus	bottle	boat	bird	bike	areo
Global Jitter (VGGNet)	77.8	67.4	92.1	91.2	85.2	56.6	92.8	92.9	90.4	97.1
SCFV (VGGNet)	79.8	66.5	92.4	91.5	86.1	59.6	90.7	93.2	90.9	96.6
CrossLayer (VGGNet)	**83.3**	**70.2**	**93.9**	**92.8**	**89.4**	59.0	**93.8**	**94.7**	**94.0**	97.8

layer. Thus it is possible to perform cross-layer pooling on two fully-connected layers. In an AConv layer, a new fully-connected layer is re-trained as one of the fully-connected layer. Certainly, it is also possible to directly use two fully-connected layers in an existing CNN without any fine-tuning, although its computational cost can be higher due to the high-dimensionality of the resulted representation, e.g., 4096×1000. Also, it is possible to perform cross-layer pooling on an original convolutional layer and a fully-connected layer in the above setting. In such a case, multiple spatial units in a convolutional layer correspond to one fully-connected layer output. To apply cross-layer pooling we can either flatten activations from multiple spatial units into a long vector or use the pooled activation from mutiple spatial units. In this paper, we use the latter since it produces lower dimensional image representations.

In this subsection, we also conduct experimental evaluation by following the above two approaches to perform cross-layer pooling. Specifically, for the first approach, denoted as FC-FC cross-layer pooling, we use the activations from fc-18 (4096 dimensions) and fc-20 (1000 dimensions) in the VGGNet model to perform cross-layer pooling. PCA is applied to reduce the dimensionality of fc-18 to 2000. For the second approach, denoted as FC-Conv cross-layer pooling, we use the activations from conv-16 and fc-16 in VGG net and apply sum-pooling (followed by using square-root post-processing) to pool the activations of conv-16. Thus the dimensionality of the representation obtained from cross-layer pooling is 512×4096.

The performance of these two methods are shown in Table 5. From Table 5, we can make the following two observations.

(a) By cross-referencing the performance in Table 1, Table 2 and Table 3, we can observe that FC-FC cross-layer pooling performs similarly to cross-layer pooling by using the AConv layer, although it tends to give slightly worse performance. This may suggest that the good performance of the AConv layer based cross-layer pooling mainly comes from the cross-layer pooling strategy although the re-trained new convolutional layer can further boost classification performance.

(b) FC-Conv cross-layer pooling achieves better performance than FC-FC cross-layer pooling on Birds200 but it is inferior to FC-FC cross-layer pooling on MIT-67 and PASCAL-07. Considering that the pooling operation of FC-Conv cross-layer pooling is performed on the OConv layer, this observation is consistent with the conclusion in Table 1 and Table 3. That is, cross-layer pooling on AConv layers leads to better performance than pooling with OConv layers for generic classification on MIT-67 and PASCAL-07.

TABLE 5
Comparison of results obtained by using cross-layer pooling with fully-connected layers.

Method	MIT-67	Birds200	PASCAL07
FC-18, FC-20 (VGGNet)	77.0%	66.1%	85.9%
Conv-15, FC-18 (VGGNet)	74.1%	68.5%	84.5%

TABLE 6
Comparison of results obtained by using different convolutional layers.

Method	MIT-67	Birds200	PASCAL07
CL-3-4 (AlexNet)	59.4%	63.9%	66.3%
CL-4-5 (AlexNet)	**63.0%**	**73.5%**	**71.3%**
CL-14-15 (VGGNet)	73.7%	**77.0%**	82.2%
CL-15-16 (VGGNet)	**74.4%**	**77.0%**	**84.1%**

4.2.2 The impact of PCA and normalization

In our implementation, we have applied three operations to obtain the final image representation, that is, peforming PCA on the local feature, performing $l2$ normalization on each pooled coding vector and power normalization. In this section, we investigate the impact of those three operations. We conduct our experiment on MIT-67 with the AConv layer features and test the performance under various settings of those three operations. Table 7 shows the results. From Table 7, we make the following observations.

(a) These three operations can have important impact on the performance. If none of them is applied, the performance drops significantly.

(b) Applying either l_2 normalization or power normalization can lead to similar performance improvement.

(c) Applying PCA with normalization, l_2 normalization or power normalization or both of them, can lead to further performance improvement.

(d) The best performance is obtained by applying those three operations altogether.

4.2.3 Feature sign quantization

As has been discussed in Section 3.4, feature sign quantization is a promising strategy to save the memory cost of cross-layer pooling. Here we demonstrate the effect of applying feature sign quantization. Feature sign quantization

TABLE 7
The impact of PCA, l_2 normalization and power normalization.

PCA	l_2 normalization	power normalization	Result
No	No	No	69.7%
No	No	Yes	73.2%
No	Yes	No	74.6%
No	Yes	Yes	72.6%
Yes	No	No	69.7%
Yes	No	Yes	77.1%
Yes	Yes	No	76.4%
Yes	Yes	Yes	**78.2%**

quantizes a feature to 1 if it is positive, -1 if it is negative and 0 if it equals to 0. In other words, we use 2 bits to represent each dimension of the pooled feature vector. This scheme greatly saves the memory usage. Here we report the result on the best performed setting for each dataset. The results are shown in Table 10.

As can be seen, this coarse quantization scheme does not degrade the performance much. On two datasets, MIT-67 and Birds-200, it achieves almost the same performance as the original feature. Note that similar quantization scheme has been also explored in [31]. However there the authors reported significant performance drop if applied to convolutional layer features. For example, in the Table 7 of [31], by binarizing conv-5, the performance drops around 5%. In contrast, our representation seems to be less sensitive to this coarse quantization.

4.3 Comparison with alternative pooling methods

Finally, we compare several alternative pooling strategies against cross-layer pooling. We first compare the baselines which directly pools convolutional layers. Moreover, the combination of multiple pooling strategies are considered in our experiment: (a) Pooling methods. We consider both sum-pooling with square-root post-processing (which is better than direct sum-pooling) and max-pooling. (b) Spatial pyramid. Three spatial pyramid partitions, that is, the 1×1 (level 0), $1 \times 1 + 2 \times 2$ (level 1) and $1 \times 1 + 2 \times 2 + 4 \times 4$ (level3) are considered. (c) Pooling layers, the conv-15 layer, conv-16 layer and the concatenation of conv-15 and conv-16 layers of VGGNet are considered. The classification results on MIT67, Birds200 and PASCAL07 are reported in Table 9. As can be seen, the best performance achieved by tuning those pooling strategies is 69.0% on MIT67, 64.1% on Birds200 and 77.4% on PASCAL07. In comparison, the cross-layer pooling counterpart achieves 74.4%, 77.0% and 84.1% on MIT67, Birds200 and PASCAL07 respectively, which is much better than the traditional pooling approach.

Another possible variation of cross-layer pooling is to change the channel pooling method from sum-pooling to max-pooling. We have also tried this setting and achieved 72.8% on MIT67, 71.6% on Birds200 and 85.2% on PASCAL07. As can be seen, its performance is slightly worse than sum-pooling. Thus we suggest using sum-pooling as the default pooling method for cross-layer pooling.

4.4 Experiments for image retrieval

For image retrieval, we evaluate cross-layer pooling on Oxford5K [32], Holiday [33] and Sculpture6K [34] datasets.

TABLE 8
Results obtained by using feature sign quantization.

Dataset	Feature sign quantization	Original
MIT-67 (VGG, Aconv)	77.9%	78.2%
Birds-200 (VGG, OConv)	76.5%	77.0%
PASCAL07 (VGG, AConv)	85.1%	87.0%

These three datasets represent several common scenarios in image retrieval.

The objects-of-interest in Oxford5K are buildings which have rich texture patterns. For the Holiday dataset, the retrieval images are more general scenes and objects. The Sculpture6K dataset focuses on sculptures which have relatively smooth surfaces. We adopt a similar setting as in [5] to evaluate performance by directly using the query image without the object bounding box.

We have re-implemented the baseline method of [5] by strictly following its experimental protocol. Specifically, we apply PCA whitening and calculate the PCA projection matrix from external datasets. For Oxford5K, we learn the PCA projection matrix from Paris6K. For Holiday, we learn the PCA projection matrix from 5K Flickr images (although those 5K Flickr images might not be same as the one used in [5]). For Sculpture6k, we also calculate the PCA projection matrix from the same 5K Flickr images.

VGGNet is used in this experiment. The baseline in [5] is performed on the conv-16 layer since it leads to the best performance. Our approach is applied on conv-15 and conv-16 layers. We use binarized cross-layer pooling vectors and select the pooling channels with top-k largest average activation value on conv-16.

Different k are evaluated and the comparison of cross-layer pooling and the baseline method in [5] is shown in Figure 4. From Figure 4, it clear that cross-layer pooling performs much better than the direct convolutional layer pooling baseline [5] once sufficient large k is chosen. Usually, when $k = 50$, a sufficient good performance can be obtained. Considering that all the features are binarized, the computational cost is still reasonable albeit the higher dimensionality than the baseline method. Note that if we directly binarized the feature obtained from the convolutional layer pooling baseline leads to significant performance drop while our method does not.

Also, it is interesting to see that when k becomes too large, e.g., getting close to using all pooling channels, the retrieval performance starts to drop. This is probably because including channels with small average activation values tends to introduce more noisy patterns during retrieval.

5 Conclusion

In this paper, we proposed a new method called cross-convolutional layer pooling to create image representations from the activations of two consective convolutional layers of a pre-trained CNN. We realize this idea on two types of implementations of convolutional layers and we show that these two different implementations may suit for different recognition tasks. Also, we propose a variation of the cross-convolutional layer pooling approach for the image retrieval

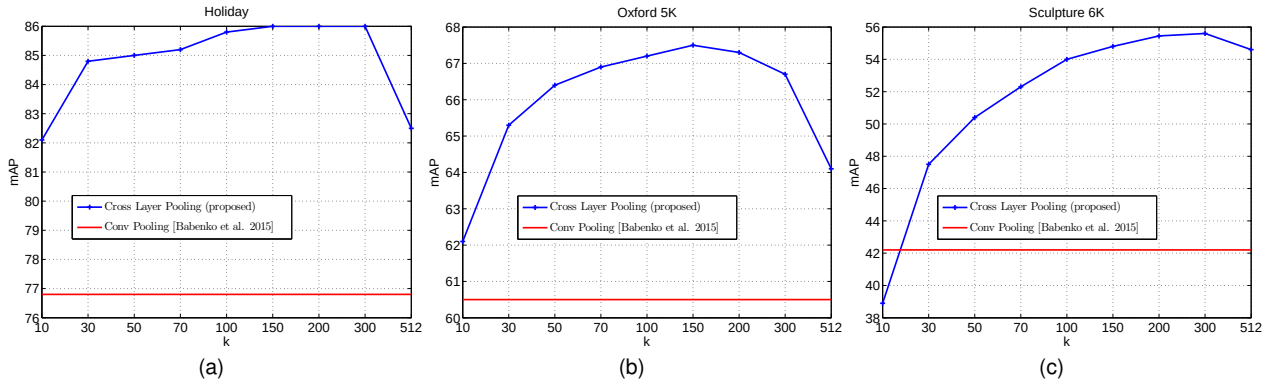

Fig. 4. The performance of cross-layer pooling on image retrieval. For our method, the feature maps of the conv-16 layer with top k maximal average activations are selected as pooling channels.

TABLE 9
Comparison of alternative pooling methods. Two pooling methods, max-pooling and sum-pooling (with square-root post-processing) are applied. Different levels of spatial pyramid configuration are applied for both pooling methods. Level 0: pooling over the whole image (1×1); level 1: 1×1+2×2; level 2: 1×1+2×2+4×4; Those pooling methods are applied to the 15th convolutional layer, the 16th convolutional layer and the concatenation of conv-15 and conv-16 layers of the VGG network. The experimental comparison is made on MIT67, Birds200 and PASCAL07 three datasets. Results are reported in the order MIT67/Birds200/PASCAL07. For reference, cross-layer pooling on conv-15 and conv-16 achieves classification accuracy 74.4/77.0/84.1, which is higher than the results of all the alternative pooling methods.

Method	Conv layer 15	Conv layer 16	Conv layer 15 + layer 16
max-pooling, SPM level0	57.3/67.7/75.5	60.7/69.0/80.5	60.6/55.8/77.8
max-pooling, SPM level1	66.3/67.6/76.8	67.7/70.2/81.6	66.0/59.1/78.9
max-pooling, SPM level2	68.0/67.6/76.8	69.3/69.9/81.9	67.9/61.5/79.8
sum-sqrt-pooling, SPM level0	66.4/66.4/77.7	68.2/69.3/81.7	69.9/65.8/80.9
sum-sqrt-pooling, SPM level1	68.0/64.1/77.5	70.0/69.8/81.7	70.7/63.6/80.6
sum-sqrt-pooling, SPM level2	69.0/64.1/77.4	70.3/69.8/81.9	70.9/64.1/80.7

TABLE 10
Results obtained by using feature sign quantization.

Dataset	Feature sign quantization	Original
MIT-67 (VGG, Aconv)	77.9%	78.2%
Birds-200 (VGG, OConv)	76.5%	77.0%
PASCAL07 (VGG, AConv)	85.1%	87.0%

task. By conducting experiments on popular image classification datasets and image retrieval datasets, we show that the proposed method leads to superior performance over various existing ways of using a pretrained DCNN to extract image representations.

ACKNOWLEDGEMENTS

This work was in part supported by the Data to Decisions Cooperative Research Centre; and ARC Grants FT120100969.

C. Shen is the corresponding author.

REFERENCES

[1] A. Krizhevsky, I. Sutskever, and G. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. Advances in Neural Inf. Process. Syst.*, 2012, pp. 1106–1114.

[2] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "ImageNet: A Large-Scale Hierarchical Image Database," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2009.

[3] A. S. Razavian, H. Azizpour, J. Sullivan, and S. Carlsson, "CNN features off-the-shelf: an astounding baseline for recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recogn. Workshop*, 2014.

[4] H. Azizpour, A. S. Razavian, J. Sullivan, A. Maki, and S. Carlsson, "From generic to specific deep representations for visual recognition," in *Proc. IEEE Conf. Computer Vision and Pattern Recogn. Workshop*, 2015.

[5] A. Babenko and V. S. Lempitsky, "Aggregating deep convolutional features for image retrieval," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015.

[6] Y. Gong, L. Wang, R. Guo, and S. Lazebnik, "Multi-scale orderless pooling of deep convolutional activation features," *Proc. Eur. Conf. Comp. Vis.*, 2014.

[7] L. Liu, C. Shen, L. Wang, A. van den Hengel, and C. Wang, "Encoding high dimensional local features by sparse coding based fisher vectors," in *Proc. Advances in Neural Inf. Process. Syst.*, 2014.

[8] Y. Li, L. Liu, C. Shen, and A. van den Hengel, "Mid-level deep pattern mining," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015.

[9] N. Zhang, R. Farrell, and T. Darrell, "Pose pooling kernels for subcategory recognition," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2012.

[10] A. v. Lingqiao Liu, Chunhua Shen, "The treasure beneath convolutional layers: Cross-convolutional-layer pooling for image classification.," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2015.

[11] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Representations*, 2015.

[12] M. Cimpoi, S. Maji, and A. Vedaldi, "Deep filter banks for texture recognition and segmentation," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, June 2015.

[13] R. B. Girshick, "Fast R-CNN," in *Proc. IEEE Int. Conf. Comp. Vis.*, 2015.

[14] T. Lin, A. RoyChowdhury, and S. Maji, "Bilinear CNN models for fine-grained visual recognition," in *Proc. IEEE Int. Conf. Comp. Vis.*, 2015.

[15] F. Perronnin, J. Sánchez, and T. Mensink, "Improving the fisher

kernel for large-scale image classification," in *Proc. Eur. Conf. Comp. Vis.*, 2010.
[16] M. Zeiler and R. Fergus, "Visualizing and understanding convolutional networks," in *Proc. Eur. Conf. Comp. Vis.*, 2014.
[17] P. W. etc., "Caltech-UCSD Birds 200," Technical Report, California Institute of Technology, 2010.
[18] L.-J. Li, H. Su, E. P. Xing, and F.-F. Li, "Object bank: A high-level image representation for scene classification and semantic feature sparsification.," in *Proc. Advances in Neural Inf. Process. Syst.*, 2010.
[19] C. Doersch, A. Gupta, and A. A. Efros, "Mid-level visual element discovery as discriminative mode seeking," in *Proc. Advances in Neural Inf. Process. Syst.*, 2013.
[20] M. Pandey and S. Lazebnik, "Scene recognition and weakly supervised object localization with deformable part-based models," in *Proc. IEEE Int. Conf. Comp. Vis.*, Washington, DC, USA, 2011, pp. 1307–1314, IEEE Computer Society.
[21] L. A. Gatys, A. S. Ecker, and M. Bethge, "Texture synthesis using convolutional neural networks," in *Proc. Advances in Neural Inf. Process. Syst.*, May 2015.
[22] A. Quattoni and A. Torralba, "Recognizing indoor scenes," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2009, pp. 413–420.
[23] M. Everingham, L. Van Gool, C. K. I. Williams, J. Winn, and A. Zisserman, "The PASCAL visual object classes challenge 2007 (VOC2007) results," http://www.pascal-network.org/challenges/VOC/voc2007/workshop/index.html.
[24] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011, Software available at http://www.csie.ntu.edu.tw/~cjlin/libsvm.
[25] N. Zhang, J. Donahue, R. Girshick, and T. Darrell, "Part-based r-cnns for fine-grained category detection," in *Proc. Eur. Conf. Comp. Vis.*, 2014.
[26] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "DeCAF: A deep convolutional activation feature for generic visual recognition," in *Proc. Int. Conf. Mach. Learn.*, 2014.
[27] N. Zhang, R. Farrell, F. Iandola, and T. Darrell, "Deformable part descriptors for fine-grained recognition and attribute prediction," in *Proc. IEEE Int. Conf. Comp. Vis.*, December 2013.
[28] Z. Song, Q. Chen, Z. Huang, Y. Hua, and S. Yan, "Contextualizing object detection and classification.," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2011.
[29] Q. Chen, Z. Song, Y. Hua, Z. Huang, and S. Yan, "Hierarchical matching with side information for image classification," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2012, pp. 3426–3433.
[30] J. Dong, W. Xia, Q. Chen, J. Feng, Z. Huang, and S. Yan, "Subcategory-aware object classification.," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2013, pp. 827–834.
[31] P. Agrawal, R. Girshick, and J. Malik, "Analyzing the performance of multilayer neural networks for object recognition," in *Proc. Eur. Conf. Comp. Vis.*, 2014.
[32] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," in *Proc. IEEE Conf. Comp. Vis. Patt. Recogn.*, 2007.
[33] M. D. Herve Jegou and C. Schmid, "Hamming embedding and weak geometry consistency for large scale image search," in *Proc. Eur. Conf. Comp. Vis.*, 2008.
[34] R. Arandjelović and A. Zisserman, "Smooth object retrieval using a bag of boundaries," in *Proc. IEEE Int. Conf. Comp. Vis.*, 2011.